massachusetts institute of technology — computer science and artificial intelligence laboratory

Asymptotics of Gaussian Regularized Least-Squares

Ross Lippert & Ryan Rifkin

AI Memo 2005-030 October 2005
CBCL Memo 257

Report Documentation Page			*Form Approved* *OMB No. 0704-0188*

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE **OCT 2005**	2. REPORT TYPE	3. DATES COVERED **00-10-2005 to 00-10-2005**
4. TITLE AND SUBTITLE **Asymptotics of Gaussian Regularized Least-Squares**		5a. CONTRACT NUMBER 5b. GRANT NUMBER 5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)		5d. PROJECT NUMBER 5e. TASK NUMBER 5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Massachusetts Institute of Technology,Artificial Intelligence Laboratory,77 Massachusetts Avenue,Cambridge,MA,02139**		8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR’S ACRONYM(S) 11. SPONSOR/MONITOR’S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES
The original document contains color images.

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**		**12**	

Asymptotics of Gaussian Regularized Least-Squares

Ross A. Lippert
M.I.T., Department of Mathematics
Building 2, Room 335
77 Massachusetts Avenue
Cambridge, MA 02139-4307
lippert@math.mit.edu

Ryan M. Rifkin
Honda Research Institute USA, Inc.
145 Tremont Street
Boston, MA 02111
rrifkin@honda-ri.com

Abstract

We consider regularized least-squares (RLS) with a Gaussian kernel. We prove that if we let the Gaussian bandwidth $\sigma \to \infty$ while letting the regularization parameter $\lambda \to 0$, the RLS solution tends to a polynomial whose order is controlled by the relative rates of decay of $\frac{1}{\sigma^2}$ and λ: if $\lambda = \sigma^{-(2k+1)}$, then, as $\sigma \to \infty$, the RLS solution tends to the kth order polynomial with minimal empirical error. We illustrate the result with an example.

1 Introduction

Given a data set $(x_1, y_1), (x_2, y_2), \ldots, (x_n, y_n)$, the inductive learning task is to build a function $f(x)$ that, given a new x point, can predict the associated y value. We study the Regularized Least-Squares (RLS) algorithm for finding f, a common and popular algorithm [2, 4] that can be used for either regression or classification:

$$\min_{f \in \mathcal{H}} \frac{1}{n} \sum_{i=1}^{n} (f(x_i) - y_i)^2 + \lambda ||f||_K^2.$$

Here, $\mathcal{H}$ is a Reproducing Kernel Hilbert Space (RKHS) [1] with associated kernel function K, $||f||_K^2$ is the squared norm in the RKHS, and λ is a regularization constant controlling the tradeoff between fitting the training set accurately and forcing smoothness of f.

[0] This report describes research done at the Center for Biological & Computational Learning, which is in the McGovern Institute for Brain Research at MIT, as well as in the Dept. of Brain & Cognitive Sciences, and which is affiliated with the Computer Sciences & Artificial Intelligence Laboratory (CSAIL).

This research was sponsored by grants from Office of Naval Research (DARPA) Contract No. MDA972-04-1-0037, Office of Naval Research (DARPA) Contract No. N00014-02-1-0915, National Science Foundation-NIH (CRCNS) Contract No. EIA-0218506, and National Institutes of Health (Conte) Contract No. 1 P20 MH66239-01A1.

Additional support was provided by Central Research Institute of Electric Power Industry (CRIEPI), Daimler-Chrysler AG, Eastman Kodak Company, Honda Research Institute USA, Inc., Komatsu Ltd., Merrill-Lynch, NEC Fund, Oxygen, Siemens Corporate Research, Inc., Sony, Sumitomo Metal Industries, and the Eugene McDermott Foundation.

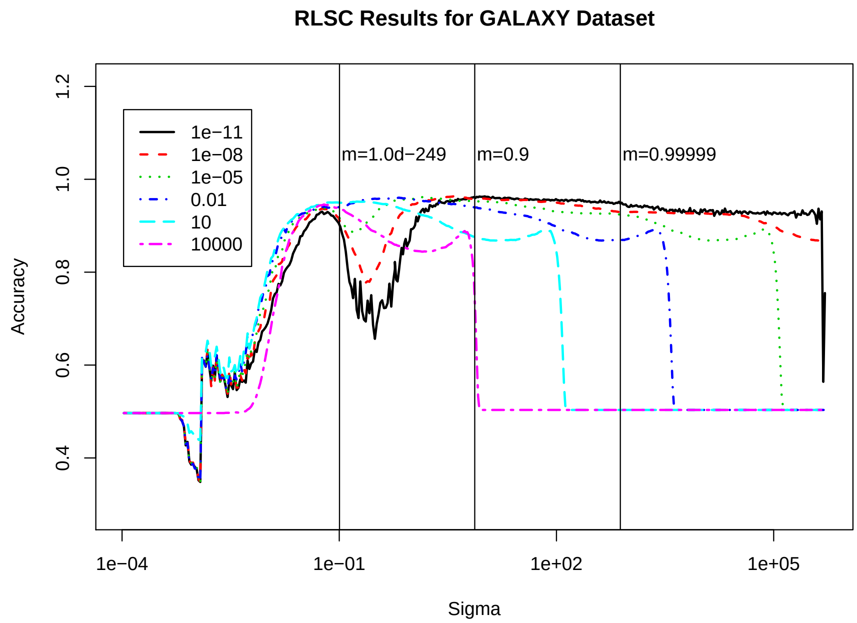

Fig. 1. RLS classification accuracy results for the UCI Galaxy dataset over a range of σ (along the x-axis) and λ (different lines) values. The vertical labelled lines show m, the smallest entry in the kernel matrix for a given σ. We see that when $\lambda = 1e-11$, we can classify quite accurately when the smallest entry of the kernel matrix is .99999.

The Representer Theorem [6] proves that the RLS solution will have the form

$$f(x) = \sum_{i=1}^{n} c_i K(x_i, x),$$

and it is easy to show [4] that we can find the coefficients c by solving the linear system

$$(K + \lambda n I)c = y, \tag{1}$$

where K is the n by n matrix satisfying $K_{ij} = K(x_i, x_j)$.

We focus on the Gaussian kernel $K(x_i, x_j) = \exp(-||x_i - x_j||^2/2\sigma^2)$.

Our work was originally motivated by the empirical observation that on a range of benchmark classification tasks, we achieved surprisingly accurate classification using a Gaussian kernel with a very large σ and a very small λ (Figure 1; additional examples in [5]). This prompted us to study the large-σ asymptotics of RLS. As $\sigma \to \infty, K(x_i, x_j) \to 1$ for arbitrary x_i and x_j. Consider a single test point x_0. RLS will first find c using Equation 1, then compute

$$f(x_0) = c^t k$$

where k is the kernel vector, $k_i = K(x_i, x_0)$. Combining the training and testing steps, we see that

$$f(x_0) = y^t (K + \lambda n I)^{-1} k$$

Both K and k are close to 1 for large σ, i.e. $K_{ij} = 1 + \epsilon_{ij}$ and $k_i = 1 + \epsilon_i$. If we directly compute $c = (K + \lambda n I)^{-1} y$, we will tend to wash out the effects of the ϵ_{ij} term as σ

becomes large. If, instead, we compute $f(x_0)$ by associating to the right, first computing *point affinities* $(K+\lambda nI)^{-1}k$, then the ϵ_{ij} and ϵ_j interact meaningfully; this interaction is crucial to our analysis.

Our approach is to Taylor expand the kernel elements (and thus K and k) in $1/\sigma$, noting that as $\sigma \to \infty$, consecutive terms in the expansion differ enormously. In computing $(K+\lambda nI)^{-1}k$, these scalings cancel each other out, and result in finite point affinities even as $\sigma \to \infty$. The asymptotic affinity formula can then be "transposed" to create an alternate expression for $f(x_0)$. Our main result is that if we set $\sigma^2 = s^2$ and $\lambda = s^{-(2k+1)}$, then, as $s \to \infty$, the RLS solution tends to the kth order polynomial with minimal empirical error.

We note in passing that our work is somewhat in the same vein as the elegant recent work of Keerthi and Lin [3]; they consider Support Vector Machines rather than RLS, and derive only the linear (first order) result.

2 Notation and definitions

Definition 1. *Let x_i be a set of $n+1$ points ($0 \le i \le n$) in a d dimensional space. The scalar x_{ia} denotes the value of the a^{th} vector component of the i^{th} point.*

The $n \times d$ matrix, X is given by $X_{ia} = x_{ia}$.

We think of X as the matrix of training data $x_1, \ldots, x_n$ and x_0 as an $1 \times d$ matrix consisting of the test point.

Let $1_m, 1_{lm}$ denote the m dimensional vector and $l \times m$ matrix with components all 1, similarly for $0_m, 0_{lm}$. We will dispense with such subscripts when the dimensions are clear from context.

Definition 2 (Hadamard products and powers). *For two $l \times m$ matrices, N, M, $N \odot M$ denotes the $l \times m$ matrix given by $(N \odot M)_{ij} = N_{ij} M_{ij}$. Analogously, we set $(N^{\odot c})_{ij} = N_{ij}^c$.*

Definition 3 (polynomials in the data). *Let $I \in \mathbb{Z}^d_{\ge 0}$ (non-negative multi-indices) and Y be a $k \times d$ matrix. Y^I is the k dimensional vector given by $\left(Y^I\right)_i = \prod_{a=1}^d Y_{ia}^{I_a}$. If $h : \mathbb{R}^d \to \mathbb{R}$ then $h(Y)$ is the k dimensional vector given by $\left(h(Y)\right)_i = h(Y_{i1}, \ldots, Y_{id})$.*

The d canonical vectors, $e_a \in \mathbb{Z}^d_{\ge 0}$, are given by $(e_a)_b = \delta_{ab}$.

For example, X^{ke_a} is the a^{th} column of X raised, elementwise, to the k^{th} power and, similarly, $x_0^{ke_a} = x_{0a}^k$. The degree of the multi-index I is $|I| = \sum_{a=1}^d I_a$. The vector $h(Y)$ where $h(y) = \sum_{a=1}^d y_a^2$ is referred to as $||Y||^2$.

In constrast, any scalar function, $f : \mathbb{R} \to \mathbb{R}$, applied to any matrix or vector, A, will be assumed to denote the elementwise application of f. We will treat $y \to e^y$ as a scalar function (we have no need of matrix exponentials in this work, so the notation is unambiguous).

We can re-express the kernel matrix and kernel vector in this notation:

$$K = e^{\frac{1}{2\sigma^2}\sum_{a=1}^d 2X^{e_a}(X^{e_a})^t - X^{2e_a}1_n^t - 1_n\left(X^{2e_a}\right)^t} \tag{2}$$

$$= \mathrm{diag}\left(e^{-\frac{1}{2\sigma^2}||X||^2}\right) e^{\frac{1}{\sigma^2}XX^t} \mathrm{diag}\left(e^{-\frac{1}{2\sigma^2}||X||^2}\right) \tag{3}$$

$$k = e^{\frac{1}{2\sigma^2}\sum_{a=1}^d 2X^{e_a}x_0^{e_a} - X^{2e_a}1_1 - 1_n x_0^{2e_a}} \tag{4}$$

$$= \mathrm{diag}\left(e^{-\frac{1}{2\sigma^2}||X||^2}\right) e^{\frac{1}{\sigma^2}Xx_0^t} e^{-\frac{1}{2\sigma^2}||x_0||^2}. \tag{5}$$

3 Orthogonal polynomial bases

Let $V_c = \text{span}\{X^I : |I| = c\}$ and $V_{\le c} = \bigcup_{a=0}^{c} V_c$ which can be thought of as the set of all d variable polynomials of degree c, evaluated on the training data. Since the data are finite, there exists b such that $V_{\le c} = V_{\le b}$ for all $c \ge b$. Generically, b is the smallest c such that $\begin{pmatrix} c+d \\ d \end{pmatrix} \ge n$.

Let Q be an orthonormal matrix in $\mathbb{R}^{n\times n}$ whose columns progressively span the $V_{\le c}$ spaces, i.e. $Q = \begin{pmatrix} B_0 & B_1 & \cdots & B_b \end{pmatrix}$ where $Q^tQ = I$ and colspan$\{\begin{pmatrix} B_0 & \cdots & B_c \end{pmatrix}\} = V_{\le c}$. We might imagine building such a Q via the Gramm-Schmidt process on the vectors $X^0, X^{e_1}, \ldots, X^{e_d}, \ldots X^I, \ldots$ taken in order of non-decreasing $|I|$.

Letting $C_I = \begin{pmatrix} |I| \\ I_1 \ldots I_d \end{pmatrix}$ be multinomial coefficients, the following relations between Q, X, and x_0 are easily proved.

$$(Xx_0^t)^{\odot c} = \sum_{|I|=c} C_I X^I (x_0^I)^t \quad \text{hence} \quad (Xx_0^t)^{\odot c} \in V_c$$

$$(XX^t)^{\odot c} = \sum_{|I|=c} C_I X^I (X^I)^t \quad \text{hence} \quad \text{colspan}\{(XX^t)^{\odot c}\} = V_c$$

and thus, $B_i^t(Xx_0^t)^{\odot c} = 0$ if $i > c$, $B_i^t(XX^t)^{\odot c}B_j = 0$ if $i > c$ or $j > c$, and $B_c^t(XX^t)^{\odot c}B_c$ is non-singular.

Finally, we note that $\text{argmin}_{v\in V_{\le c}}\{||y - v||\} = \sum_{a\le c} B_a(B_a^t y)$.

4 Taking the $\sigma \to \infty$ limit

We will begin with a few simple lemmas about the limiting solutions of linear systems. At the end of this section we will arrive at the limiting form of suitably modified RLSC equations.

Lemma 1. *Let $A(s)$ be a continuous matrix-valued function defined for $0 < s < s_0$ for some $s_0 \in \mathbb{R}$. If $\lim_{s\to 0} A(s) = A_0$ and A_0 is non-singular, then $\lim_{s\to 0} A(s)^{-1} = A_0^{-1}$.*

Proof. Given ϵ, select $\delta < s_0$ such that $||I - A(s)A_0^{-1}||_2 < \min\left\{\frac{1}{2}, \frac{\epsilon}{2||A_0^{-1}||_2}\right\}$ for $s < \delta$ (such a δ exists since $\lim_{s\to 0} A(s) = A_0$). Note that $||I - A(s)A_0^{-1}||_2 < \frac{1}{2}$, implies $A(s)$ is non-singular. Then

$$A(s)^{-1} = A_0^{-1}(I - (I - A(s)A_0^{-1}))^{-1} = A_0^{-1}\left(I + \sum_{i\ge 1}(I - A(s)A_0^{-1})^i\right)$$

$$||A_0^{-1} - A(s)^{-1}||_2 \le ||A_0^{-1}||_2 \frac{||I - A(s)A_0^{-1}||_2}{1 - ||I - A(s)A_0^{-1}||_2} < \epsilon.$$

□

Corollary 1. *Let $A(s), y(s)$ be continuous matrix-valued and vector-valued functions, defined for $0 < s < s_0$ for some $s_0 \in \mathbb{R}$ with $\lim_{s\to 0} A(s) = A_0$ is non-singular. $\lim_{s\to 0} y(s) = y_0$ iff $\lim_{s\to 0} A(s)^{-1}y(s) = A_0^{-1}y_0$.*

Proof. By lemma 1, $\lim_{s\to 0} A(s)^{-1} = A_0^{-1}$.

By the continuity of matrix multiplication

$$\lim_{s\to 0} B(s)x(s) = \left(\lim_{s\to 0} B(s)\right)\left(\lim_{s\to 0} x(s)\right)$$

(the existence of the right hand limits implying the existence of the left hand limit).

If $\lim_{s\to 0} y(s) = y_0$ then let $B(s) = A^{-1}(s)$ and $x(s) = y(x)$.

If $\lim_{s\to 0} A(s)^{-1}y(s) = x_0$ then let $x(s) = A(s)^{-1}y(s)$ and $B(s) = A(s)$, and thus $y_0 = \lim_{s\to 0} A(s)(A(s)^{-1}y(s)) = A_0x_0$. $\square$

Lemma 2. *Let $A(s), y(s)$ be matrix-valued and vector-valued polynomials of degree p and $B(s), z(s)$ be matrix-valued and vector-valued functions that are bounded in the region $0 < s < s_0$, for some $s_0 \in \mathbb{R}$. If $A(s)$ is non-singular for $0 < s < s_0$, then*

$$\lim_{s\to 0}(A(s) + s^{p+1}B(s))^{-1}(y(s) + s^{p+1}z(s)) = \lim_{s\to 0} A(s)^{-1}y(s).$$

Proof. We first note that for $s > 0$,

$$(A(s) + s^{p+1}B(s))^{-1} = (I + s^{p+1}A(s)^{-1}B(s))^{-1}A(s)^{-1}$$

Since $A(s)$ is a polynomial, the entries of $A(s)^{-1}$ are rational functions with denominators of degree p. Thus, $\lim_{s\to 0} s^{p+1}A^{-1}(s) = 0$, and thus, by the boundedness of $B(s)$ and $z(s)$,

$$\begin{aligned} s^{p+1}A^{-1}(s)z(s) &\to 0 \\ s^{p+1}A^{-1}(s)B(s) &\to 0. \end{aligned}$$

By Lemma 1, $\lim_{s\to 0}(I + s^{p+1}A^{-1}(s)B(s)) = I$. Thus, by Corollary 1,

$$\begin{aligned} &\lim_{s\to 0}(A(s) + s^{p+1}B(s))^{-1}(y(s) + s^{p+1}z(s)) \\ &= \lim_{s\to 0}(I + s^{p+1}A(s)^{-1}B(s))^{-1}A(s)^{-1}(y(s) + s^{p+1}z(s)) \\ &= \lim_{s\to 0} A(s)^{-1}(y(s) + s^{p+1}z(s)) \\ &= \lim_{s\to 0} A(s)^{-1}y(s). \end{aligned}$$

$\square$

Lemma 3. *Let $i_1 < \cdots < i_q$ be positive integers. Let $A(s), y(s)$ be a block matrix and block vector given by*

$$A(s) = \begin{pmatrix} A_{00}(s) & s^{i_1}A_{01}(s) & \cdots & s^{i_q}A_{0q}(s) \\ s^{i_1}A_{10}(s) & s^{i_1}A_{11}(s) & \cdots & s^{i_q}A_{1q}(s) \\ \cdots & \cdots & \cdots & \cdots \\ s^{i_q}A_{q0}(s) & s^{i_q}A_{q1}(s) & \cdots & s^{i_q}A_{qq}(s) \end{pmatrix}, \quad y(s) = \begin{pmatrix} b_0(s) \\ s^{i_1}b_1(s) \\ \cdots \\ s^{i_q}b_q(s) \end{pmatrix}$$

where $A_{ij}(s)$ and $b_i(s)$ are continuous matrix-valued and vector-valued functions of s with $A_{ii}(0)$ non-singular for all i.

$$\lim_{s\to 0} A^{-1}(s)y(s) = \begin{pmatrix} A_{00}(0) & 0 & \cdots & 0 \\ A_{10}(0) & A_{11}(0) & \cdots & 0 \\ \cdots & \cdots & \cdots & \cdots \\ A_{q0}(0) & A_{q1}(0) & \cdots & A_{qq}(0) \end{pmatrix}^{-1} \begin{pmatrix} b_0(0) \\ b_1(0) \\ \cdots \\ b_q(0) \end{pmatrix}$$

Proof. Let $P(s) = \text{diag}(I, s^{-i_1}I, \dots, s^{-i_q}I)$ with the blocks of $P(s)$ commensurate with those of $A(s)$.

$$P(s)A(s) = \begin{pmatrix} A_{00}(s) & s^{i_1}A_{01}(s) & \cdots & s^{i_q}A_{0q}(s) \\ A_{10}(s) & A_{11}(s) & \cdots & s^{i_q-i_1}A_{1q}(s) \\ \cdots & \cdots & \cdots & \cdots \\ A_{q0}(s) & A_{q1}(s) & \cdots & A_{qq}(s) \end{pmatrix}$$

and

$$\lim_{s\to 0} P(s)A(s) = \begin{pmatrix} A_{00}(0) & 0 & \cdots & 0 \\ A_{10}(0) & A_{11}(0) & \cdots & 0 \\ \cdots & \cdots & \cdots & \cdots \\ A_{q0}(0) & A_{q1}(0) & \cdots & A_{qq}(0) \end{pmatrix}^{-1}$$

which is invertible.

Noting that $\lim_{s\to 0} P(s)y(s) = \begin{pmatrix} b_0(s) \\ b_1(s) \\ \cdots \\ b_q(s) \end{pmatrix}$, we see that our result follows from corollary 1 applied to $\lim_{s\to 0}(P(s)A(s))^{-1}(P(s)y(s))$. $\square$

We are now ready to state and prove the main result of this section, characterizing the limiting large-σ solution of Gaussian RLS.

Theorem 1. *Let q be an integer satisfying $q < b$, and let $p = 2q + 1$. Let $\lambda = C\sigma^{-p}$ for some constant C. Define $A_{ij}^{(c)} = \frac{1}{c!}B_i^t(XX^t)^{\odot c}B_j$, and $b_i^{(c)} = \frac{1}{c!}B_i^t\,(Xx_0^t)^{\odot c}$.*

$$\lim_{\sigma\to\infty}\left(K + nC\sigma^{-p}I\right)^{-1}k = v$$

where

$$v = \begin{pmatrix} B_0 & \cdots & B_q \end{pmatrix} w \tag{6}$$

$$\begin{pmatrix} b_0^{(0)} \\ b_1^{(1)} \\ \cdots \\ b_q^{(q)} \end{pmatrix} = \begin{pmatrix} A_{00}^{(0)} & 0 & \cdots & 0 \\ A_{10}^{(1)} & A_{11}^{(1)} & \cdots & 0 \\ \cdots & \cdots & \cdots & \cdots \\ A_{q0}^{(q)} & A_{q1}^{(q)} & \cdots & A_{qq}^{(q)} \end{pmatrix} w \tag{7}$$

We first manipulate the equation $(K + n\lambda I)y = k$ according to the factorizations in (3) and (5). Defining

$$N \equiv \text{diag}\, e^{-\frac{1}{2\sigma^2}||X||^2}, \quad \alpha \equiv e^{-\frac{1}{2\sigma^2}||x_0||^2},$$
$$P \equiv e^{\frac{1}{\sigma^2}XX^t}, \quad w \equiv e^{\frac{1}{\sigma^2}Xx_0^t}, \quad \beta \equiv nC\sigma^{-p},$$

(where we omit for brevity the dependencies on σ) we have

$$K = \text{diag}\left(e^{-\frac{1}{2\sigma^2}||X||^2}\right)e^{\frac{1}{\sigma^2}XX^t}\text{diag}\left(e^{-\frac{1}{2\sigma^2}||X||^2}\right) = NPN$$
$$k = \text{diag}\left(e^{-\frac{1}{2\sigma^2}||X||^2}\right)e^{\frac{1}{\sigma^2}Xx_0^t}e^{-\frac{1}{2\sigma^2}||x_0||^2} = Nw\alpha$$

Noting that

$$\lim_{\sigma\to\infty} e^{-\frac{1}{2\sigma^2}||x_0||^2}\text{diag}\left(e^{\frac{1}{2\sigma^2}||X||^2}\right) = \lim_{\sigma\to\infty}\alpha N^{-1} = I,$$

we have

$$\begin{aligned} v &\equiv \lim_{\sigma\to\infty}(K + nC\sigma^{-p}I)^{-1}k \\ &= \lim_{\sigma\to\infty}(NPN + \beta I)^{-1}Nw\alpha \\ &= \lim_{\sigma\to\infty}\alpha N^{-1}(P + \beta N^{-2})^{-1}w \\ &= \lim_{\sigma\to\infty}\alpha N^{-1}(P + \beta N^{-2})^{-1}w \\ &= \lim_{\sigma\to\infty}\left(e^{\frac{1}{\sigma^2}XX^t} + nC\sigma^{-p}\text{diag}\left(e^{\frac{1}{\sigma^2}||X||^2}\right)\right)^{-1} e^{\frac{1}{\sigma^2}Xx_0^t}. \end{aligned}$$

Changing bases with Q,

$$Q^t v = \lim_{\sigma\to\infty}\left(Q^t e^{\frac{1}{\sigma^2}XX^t}Q + nC\sigma^{-p}Q^t\text{diag}\left(e^{\frac{1}{\sigma^2}||X||^2}\right)Q\right)^{-1} Q^t e^{\frac{1}{\sigma^2}Xx_0^t}.$$

Expanding via Taylor series and writing in block form (in the $b \times b$ block structure of Q),

$$\begin{aligned} Q^t e^{\frac{1}{\sigma^2}XX^t}Q &= Q^t(XX^t)^{\odot 0}Q + \frac{1}{1!\sigma^2}Q^t(XX^t)^{\odot 1}Q + \frac{1}{2!\sigma^4}Q^t(XX^t)^{\odot 2}Q + \cdots \\ &= \begin{pmatrix} A_{00}^{(0)} & 0 & \cdots & 0 \\ 0 & 0 & \cdots & 0 \\ \cdots & \cdots & \cdots & \cdots \\ 0 & 0 & \cdots & 0 \end{pmatrix} + \frac{1}{\sigma^2}\begin{pmatrix} A_{00}^{(1)} & A_{01}^{(1)} & \cdots & 0 \\ A_{10}^{(1)} & A_{11}^{(1)} & \cdots & 0 \\ \cdots & \cdots & \cdots & \cdots \\ 0 & 0 & \cdots & 0 \end{pmatrix} + \cdots \end{aligned}$$

$$\begin{aligned} Q^t e^{\frac{1}{\sigma^2}Xx_0^t} &= Q^t(Xx_0^t)^{\odot 0} + \frac{1}{\sigma^2}Q^t(Xx_0^t)^{\odot 1} + \frac{1}{\sigma^4}Q^t(Xx_0^t)^{\odot 2} + \cdots \\ &= \begin{pmatrix} b_0^{(0)} \\ 0 \\ \cdots \\ 0 \end{pmatrix} + \frac{1}{\sigma^2}\begin{pmatrix} b_0^{(1)} \\ b_1^{(1)} \\ \cdots \\ 0 \end{pmatrix} + \cdots \end{aligned}$$

$$nC\sigma^{-p}Q^t\text{diag}\left(e^{\frac{1}{\sigma^2}||X||^2}\right)Q = nC\sigma^{-p}I + \cdots.$$

Since the $A_{cc}^{(c)}$ are non-singular, Lemma 3 applies, giving our result. □

5 The classification function

When performing RLS, the actual prediction of the limiting classifier is given via

$$f_\infty(x_0) \equiv \lim_{\sigma\to\infty} y^t(K + nC\sigma^{-p}I)^{-1}k.$$

Theorem 1 determines

$$v = \lim_{\sigma\to\infty}(K + nC\sigma^{-p}I)^{-1}k,$$

showing that $f_\infty(x_0)$ is a polynomial in the training data X. In this section, we show that $f_\infty(x_0)$ is, in fact, a polynomial in the test point x_0. We continue to work with the orthonormal vectors B_i as well as the auxilliary quantities $A_{ij}^{(c)}$ and $b_i^{(c)}$ from Theorem 1.

Theorem 1 shows that $v \in V_{\leq q}$: the point affinity function is a polynomial of degree q in the training data, determined by (7).

$$\sum_{i,j \leq c} c! B_i A_{ij}^{(c)} B_j^t = (XX^t)^{\odot c} \quad \text{hence} \quad \sum_{j \leq c} c! B_c A_{cj}^{(c)} B_j^t = B_c B_c^t (XX^t)^{\odot c}$$
$$\sum_{i \leq c} c! B_i b_i^{(c)} = (X x_0^t)^{\odot c} \quad \text{hence} \quad c! B_c b_i^{(c)} = B_c B_c^t (X x_0^t)^{\odot c}$$

we can restate Equation 7 in an equivalent form:

$$\begin{pmatrix} B_0^t \\ \cdots \\ B_q^t \end{pmatrix}^t \left(\begin{pmatrix} 0! b_0^{(0)} \\ 1! b_1^{(1)} \\ \cdots \\ q! b_q^{(q)} \end{pmatrix} - \begin{pmatrix} 0! A_{00}^{(0)} & 0 & \cdots & 0 \\ 1! A_{10}^{(1)} & 1! A_{11}^{(1)} & \cdots & 0 \\ \cdots & \cdots & \cdots & \cdots \\ q! A_{q0}^{(q)} & q! A_{q1}^{(q)} & \cdots & q! A_{qq}^{(q)} \end{pmatrix} \begin{pmatrix} B_0^t \\ \cdots \\ B_q^t \end{pmatrix} v \right) = 0 \tag{8}$$
$$\sum_{c \leq q} c! B_c b_c^{(c)} - \sum_{c \leq q} \sum_{j \leq c} c! B_c A_{cj}^{(c)} B_j^t v = 0 \tag{9}$$
$$\sum_{c \leq q} B_c B_c^t \left((X x_0^t)^{\odot c} - (XX^t)^{\odot c} v \right) = 0. \tag{10}$$

Up to this point, our results hold for arbitrary training data X. To proceed, we require a mild condition on our training set.

Definition 4. *X is called* generic *if $X^{I_1}, \ldots, X^{I_n}$ are linearly independent for any distinct multi-indices $\{I_i\}$.*

Lemma 4. *For generic X, the solution to Equation 7 (or equivalently, Equation 10) is determined by the conditions*

$$\forall I : |I| \leq q, (X^I)^t v = x_0^I, \tag{11}$$

where $v \in V_{\leq q}$.

Proof. By definition, $V_{\leq q} = \text{span}\{X^I : |I| \leq q\}$ and, by genericity, the vectors X^I where $|I| \leq q < b$ are linearly independent. Thus (11) reduces to a $\begin{pmatrix} q+d \\ d \end{pmatrix} \times \begin{pmatrix} q+d \\ d \end{pmatrix}$ system of linear equations with unique solution, which we will call v. We now show that v satisfies (10).

$$(XX^t)^{\odot c} = \sum_{|I|=c} C_I X^I (X^I)^t \quad \text{and} \quad (X x_0^t)^{\odot c} = \sum_{|I|=c} C_I X^I (x_0^I)^t$$
$$\sum_{|I|=c} C_I X^I (X^I)^t v \;=\; \sum_{|I|=c} C_I X^I x_0^I.$$

and thus $(XX^t)^{\odot c} v = (X x_0^t)^{\odot c}$. ☐

Theorem 2. *For generic data, let v be the solution to Equation 10. For any $y \in \mathbb{R}^n$, $f(x_0) = y^t v = h(x_0)$, where $h(x) = \sum_{|I| \leq q} a_I x^I$ is a multivariate polynomial of degree q minimizing $||y - h(X)||$.*

Proof. Since $h(X)$ is the minimizer of $||y - h(X)||$,

$$h(X) = \begin{pmatrix} B_0 & \cdots & B_q \end{pmatrix} \begin{pmatrix} B_0 & \cdots & B_q \end{pmatrix}^t y.$$

Thus,

$$h(X)^t v = y^t \begin{pmatrix} B_0 & \cdots & B_q \end{pmatrix} \begin{pmatrix} B_0 & \cdots & B_q \end{pmatrix}^t v = y^t v$$

since $v \in V_{\leq q}$.

By Lemma 5,

$$h(X)^t v = \sum_{|I| \leq q} a_I (X^I)^t v = \sum_{|I| \leq q} a_I x_0^I = h(x_0).$$

□

We see that as $\sigma \to \infty$, the RLS solution tends to the minimum empirical error kth order polynomial.

6 Experimental Verification

In this section, we present a simple experiment that illustrates our results. We consider the fifth-degree polynomial function

$$f(x) = .5(1-x) + 150x(x-.25)(x-.3)(x-.75)(x-.95),$$

over the range $x \in [0, 1]$. Figure 2 plots f, along with a 150 point dataset drawn by choosing x_i uniformly in $[0, 1]$, and choosing $y = f(x) + \epsilon_i$, where ϵ_i is a Gaussian random variable with mean 0 and standard deviation .05. Figure 2 also shows (in red) the best polynomial approximations to the data (not to the ideal f) of various orders. (We omit third order because it is nearly indistinguishable from second order.)

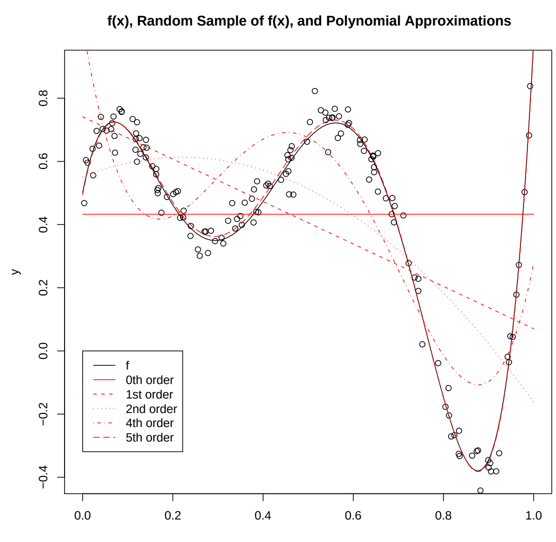

Fig. 2. $f(x) = .5(1-x) + 150x(x-.25)(x-.3)(x-.75)(x-.95)$, a random dataset drawn from $f(x)$ with added Gaussian noise, and data-based polynomial approximations to f.

According to Corollary 1, if we parametrize our system by a variable s, and solve a Gaussian regularized least squares problem with $\sigma^2 = s^2$ and $\lambda = Cs^{-(2k+1)}$ for some integer

k, then, as $s \to \infty$, we expect the solution to the system to tend to the kth-order data-based polynomial approximation to f. Asymptotically, the value of the constant C does not matter, so we (arbitrarily) set it to be 1. Figure 3 demonstrates this result.

We note that these experiments frequently require setting λ much smaller than machine-ϵ. As a consequence, we need more precision than IEEE double-precision floating-point, and our results cannot be obtained via many standard tools (e.g., MATLAB(TM)) We performed our experiments using CLISP, an implementation of Common Lisp that includes arithmetic operations on arbitrary-precision floating point numbers.

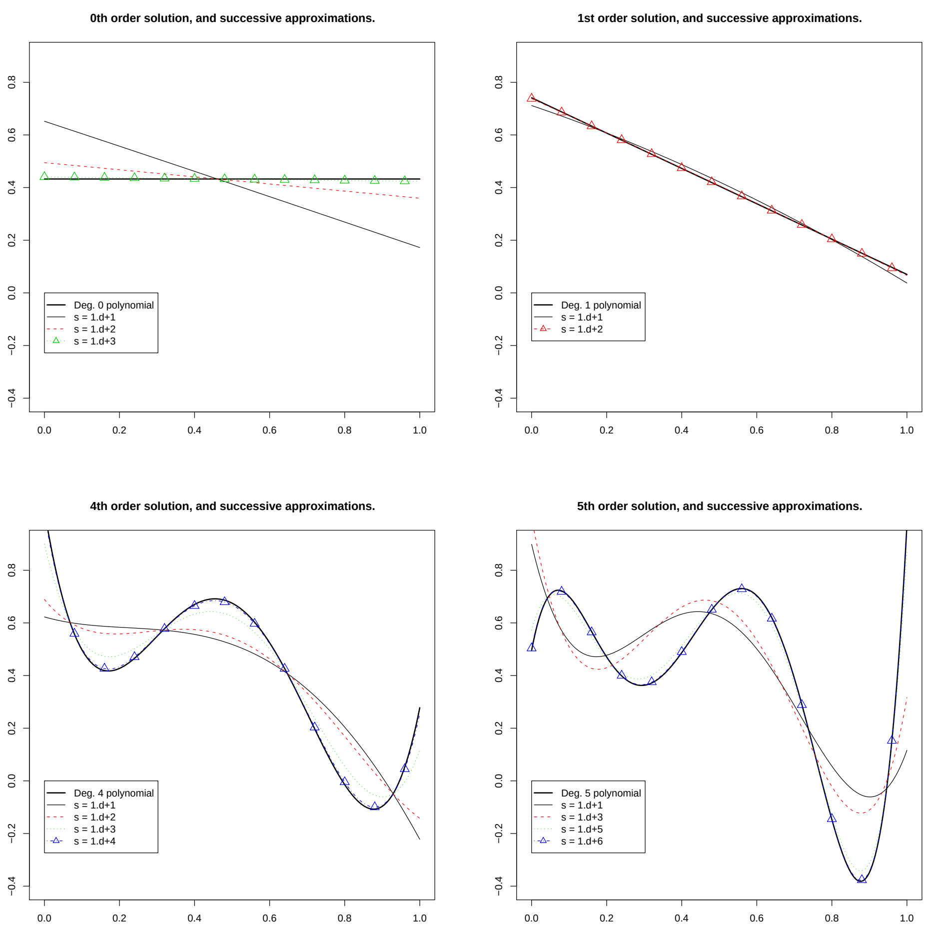

Fig. 3. As $s \to \infty$, $\sigma^2 = s^2$ and $\lambda = s^{-(2k+1)}$, the solution to Gaussian RLS approaches the kth order polynomial solution.

7 Discussion

Our result provides insight into the asymptotic behavior of RLS, and (partially) explains Figure 1: in conjunction with additional experiments not reported here, we believe that we are recovering second-order polynomial behavior, with the drop-off in performance at various λ's occurring at the transition to third-order behavior, which cannot be accurately recovered in IEEE double-precision floating-point. Although we used the specific details of RLS in deriving our solution, we expect that in practice, a similar result would hold for Support Vector Machines, and perhaps for Tikhonov regularization with convex loss more generally.

An interesting implication of our theorem is that for very large σ, we can obtain various order polynomial classifications by sweeping λ. In [5], we present an algorithm for solving for a wide range of λ for essentially the same cost as using a single λ. This algorithm is not currently practical for large σ, due to the need for extended-precision floating point.

Our work also has implications for approximations to the Gaussian kernel. Yang et al. use the Fast Gauss Transform (FGT) to speed up matrix-vector multiplications when performing RLS [7]. In [5], we studied this work; we found that while Yang et al. used moderate-to-small values of σ (and did not tune λ), the FGT sacrificed substantial accuracy compared to the best achievable results on their datasets. We showed empirically that the FGT becomes much more accurate at larger values of σ; however, at large-σ, it seems likely we are merely recovering low-order polynomial behavior. We suggest that approximations to the Gaussian kernel must be checked carefully, to show that they produce sufficiently good results are moderate values of σ; this is a topic for future work.

References

1. Aronszajn. Theory of reproducing kernels. *Transactions of the American Mathematical Society*, 68:337–404, 1950.
2. Evgeniou, Pontil, and Poggio. Regularization networks and support vector machines. *Advances In Computational Mathematics*, 13(1):1–50, 2000.
3. Keerthi and Lin. Asymptotic behaviors of support vector machines with gaussian kernel. *Neural Computation*, 15(7):1667–1689, 2003.
4. Rifkin. *Everything Old Is New Again: A Fresh Look at Historical Approaches to Machine Learning*. PhD thesis, Massachusetts Institute of Technology, 2002.
5. Rifkin and Lippert. Practical regularized least-squares: λ-selection and fast leave-one-out-computation. In preparation, 2005.
6. Wahba. *Spline Models for Observational Data*, volume 59 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. Society for Industrial & Applied Mathematics, 1990.
7. Yang, Duraiswami, and Davis. Efficient kernel machines using the improved fast Gauss transform. In *Advances in Neural Information Processing Systems*, volume 16, 2004.